

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re PATENT APPLICATION OF:

SEED, Steven et al.

Application Serial No.: **10/073,938**

Application Filing Date: 02/14/2002

Title: **MANAGED OBJECT REPLICATION AND
DELIVERY**

Attorney Docket: 0092-US-01

Group Art Unit: 2443

Examiner: BILGRAMI, Asghar

Confirmation No.: 5330

Date: November 29, 2011

APPEAL BRIEF UNDER 37 CFR § 41.37

via EFS-Web

Hon. Commissioner of Patents

P.O. Box 1450

Alexandria, VA 22313-1450

Sir:

Applicant hereby submits an Appeal Brief under 37 CFR § 41.37, following the Notice of Appeal for this application on August 31, 2011. The fee set forth in § 41.20(b)(2) is being submitted herewith. If for any reason the payment herewith is missing or insufficient or additional fees are required, the Commissioner is hereby authorized to charge the insufficiency to Deposit Account No. 50-5063.

TABLE OF CONTENTS

41.37(C)(1)(I) REAL PARTY IN INTEREST.....	3
41.37(C)(1) (II) RELATED APPEALS AND INTERFERENCES	3
41.37(C)(1) (III) STATUS OF CLAIMS	3
41.37(C)(1) (IV) STATUS OF AMENDMENTS	4
41.37(C)(1) (V) SUMMARY OF CLAIMED SUBJECT MATTER	4
<i>Claim 1</i>	5
<i>Claim 16</i>	7
<i>Claim 23</i>	8
<i>Claim 38</i>	9
<i>Claim 45</i>	10
41.37(C)(1) (VI) GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL.....	12
41.37(C)(1)(VII) ARGUMENT	12
<i>Claim 1 and its dependents are patentable over Jungck and Sim</i>	14
<i>Claim 16 and its dependents are patentable over Jungck and Sim</i>	26
<i>Claim 23 and its dependents are patentable over Jungck and Sim</i>	27
<i>Claim 38 and its dependents are patentable over Jungck and Sim</i>	28
<i>Claim 45 and its dependents are patentable over Jungck and Sim</i>	29
41.37(C)(1) (VIII) CLAIMS APPENDIX.....	30
41.37(C)(1)(IX) EVIDENCE APPENDIX.....	51
41.37(C)(1)(X) RELATED PROCEEDINGS APPENDIX	51
CONCLUSION.....	51

41.37(C)(1)(I) REAL PARTY IN INTEREST

The real party in interest is *Level 3 Communications, LLC* of Broomfield, Colorado, which is the assignee of the present application.

41.37(C)(1) (II) RELATED APPEALS AND INTERFERENCES

U.S. Patent Application No. 11/980,672 (“the ‘672 Application”) is a continuation of the present application. The ‘672 Application is on appeal, awaiting a decision by the Board of Appeals.

On information and belief, there are no other appeals, interferences, or judicial proceedings known to the appellant, the appellant's legal representative, or assignee which may be related to, directly affect or be directly affected by or have a bearing on the Board of Patent Appeals and Interferences (the "Board") decision in the pending appeal.

41.37(C)(1) (III) STATUS OF CLAIMS

Pending: Claims 1-3, 5-12, 14-18, 20-25, 27-34, 36-40, 42-47, 49-55, 57, 58, and 66-84 are pending.

Canceled: Claims 4, 13, 19, 26, 35, 41, 48, 56, and 59-65 were canceled.

Appealed: Claims 1-3, 5-12, 14-18, 20-25, 27-34, 36-40, 42-47, 49-55, 57, 58, and 66-84 (all pending claims) are appealed.

Claims 1-3, 5-12, 14-18, 20-25, 27-34, 36-40, 42-47, 49-55, 57, 58, and 66-84 (all pending claims) stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Jungck (U.S. Pub. No. 2005/0021863 – hereinafter “Jungck”) and Sim (U.S. Pub. No. 2003/0031176 – hereinafter “Sim”).

41.37(C)(1) (IV) STATUS OF AMENDMENTS

An Amendment for this application was filed on June 28, 2011, which resulted in the final Office Action dated August 3, 2011 (hereinafter the “Final Office Action”). No amendments have been filed after the final rejection.

41.37(C)(1) (V) SUMMARY OF CLAIMED SUBJECT MATTER

The **independent claims** are each discussed individually below, but a preface is given here to assist the Board in understanding the context of the invention.

This invention relates to content delivery and, more particularly, to managed object replication in a network such as a content delivery network (CDN) *Specification*, pgs. 1, 2, ¶¶0001, 0010. A typical CDN deploys one or more parent servers, hosting a plurality of objects, in a network and one or more edge servers at the edge of the network to facilitate more cost-effective and efficient delivery of such objects to an end-user (client). *Id.*, pg. 2, ¶0010. It is typically desirable to serve objects from edge servers because the edge servers are typically closer (by various measures of distance) to end-users. *Id.* For example, streaming content data from edge servers saves parent-to-edge bandwidth. *Id.* However, it is typically not feasible to store all objects on the edge servers. *Id.*, pg. 2, ¶0011. The main difficulty is due to the fact that many such objects are very large (typically on the order of 10 MB – in the neighborhood of 500 MB for movies). The storage required to accommodate often large and sometimes rarely requested objects at every edge server can be cost prohibitive as the number of customers grows and the number of their objects increases. *Id.* One solution to this problem is to pre-populate edge servers with objects for which there will likely be a significant or high demand. *Id.*, pg. 3, ¶0012. However, as the inventors noted, it is difficult to predict popularity and difficult to manage pre-populating. *Id.* Other solutions

fetch objects on demand. *Id.*, pg. 3, ¶0013. In such schemes, when a requested object is not available on a handling edge server, a connection is made between a parent server having the requested object and the handling edge server to fetch the requested object from the parent server. *Id.* Such fetching suffers however from having to go through a network path between the handling edge server and the parent server with the object whenever a client requests an object that is not already at the particular edge server. *Id.* Fetching a large object to the handling edge server through a parent path can be slow, e.g., because there may be limited available bandwidth from the parent server to the handling edge server. *Id.*, pg. 4, ¶0014. If a parent server uses too much bandwidth copying an object to an edge server, this can create congestion at that parent server. *Id.* Other problems with this approach are set forth in the application as filed, e.g. *Id.*, pg. 4, ¶0014-0015. The inventors thus realized that it would be advantageous to populate edge servers with the most popular objects yet somehow serve the rest from parent servers with a goal to maximize the amount of object bits served from edge servers of the network. *Id.*, pg. 4, ¶0016. The inventors also realized that it would also be advantageous to populate edge servers by, for example, storage fill on demand when an object is popular enough, without having to make the end-user wait for such population. *Id.*

In addition, the inventors realized, *inter alia*, that object replication should, in some cases, be a function of a measure of (I) popularity of an object *and* (II) a dynamic popularity threshold. *Id.*, pg. 17, ¶0042-43.

Claim 1

As recited in independent claim 1, the invention described and claimed in the present application relates to a method for managed object replication and delivery in a system comprising a network having one or more edge server sites

and one or more parent server sites distinct from the edge server sites. See *Specification*, Fig. 1 and pgs. 6-7, ¶¶0020-0022. As recited in claim 1, the method comprises directing a request by a client for an object to a particular edge server site in the network. *Id.*, pg. 5, ¶0017 and pg. 8, ¶0024, at 200 (in Fig. 2) and 300 (in Fig. 4).

As recited in claim 1, the method further recites: “if the particular edge server site does not have the requested object, then (b1) said particular edge server site redirecting the client request to a first server site distinct from the edge server site.” *Id.*, pgs. 9-10, ¶¶0026-0028 and Figs. 2 and 3 at 210, 310.

In addition to directing the client to another site (“a first server site”) when the first edge server does not have the requested object, claim 1 further recites “conditionally replicating a portion of the requested object on the particular edge server site for use in serving future client requests.” *Id.*, pg. 9, ¶¶0026-0027 and Figs. 2 and 3 at 315; pg. 26, ¶0057. This conditional replication is made, as recited in the claim, “based at least in part on a dynamic measure of popularity of the requested object.” (See also *Id.*, pg. 12, ¶0032)

The claim further recites that “the portion of the requested object is replicated on the particular edge server site when

- [i]* the dynamic measure of popularity of the requested object
exceeds
- [ii]* a dynamic replication threshold.”¹

Id., pgs. 16-17, ¶¶0041-43.

The claim additionally recites that the dynamic replication threshold is

¹ The labels “[i]” and “[ii]” are used here to merely to aid explanation, and are not intended to impose any limitations on the claim.

“based at least in part on at least one dynamic measure of capacity.” The “measure of popularity of the requested object” (labeled [i] for this discussion) is discussed in the Specification, e.g., under the section titled “Determining Popularity” on pgs. 14-16, ¶¶0035-0040 and Fig. 4. The “dynamic popularity threshold” (labeled [ii] for this discussion) is discussed in the Specification, e.g., at pgs. 16-17, ¶¶0041-43, under the heading “Initiating Replication.”

Claim 1 further recites that “if the particular edge server site does not have the requested object, then ... conditionally replicating a portion of the requested object on the particular edge server site for use in serving future client requests, ... wherein *the requested object is served to the client from a server site other than the particular edge server site.*” *Id.* pg. 10, ¶0028, Fig. 2, step 215, Fig. 3(a), step 345. In other words, as recited in claim 1, the requested object is selectively copied to the edge server even though it is not served from the edge server.

Claim 16

As recited in independent claim 16 the invention described and claimed in the present application recites a method for managed object replication and delivery in a system comprising a network having one or more parent server sites and one or more edge server sites distinct from the parent server sites. *Specification*, Fig. 1 and pgs. 6-7, ¶¶0020-0022. The method of claim 16 includes “directing a request by a client for an object to an optimal edge server site in the network.” *Id.*, pg. 5, ¶0017 and pg. 8, ¶0024, at 200 (in Fig. 2) and 300 (in Fig. 4). The method of claim 16 further recites “if the edge server site does not have the requested object, then the edge server site redirecting the client request to a parent server site in the network distinct from the edge server site.” *Id.*, pgs. 9-10, ¶¶0026-0028 and Figs. 2 and 3 at 210, 310.

The method of claim 16 further recites: “if the edge server site does not

have the requested object, conditionally replicating a portion of the requested object to the edge server site from a parent server site in the network, said replicating being based at least in part on

[a] a dynamic measure of popularity of the requested object
relative to

[b] a dynamic replication threshold.”²

Id., pg. 9, ¶¶0026-0027 and Figs. 2 and 3 at 315; pgs. 14-16, ¶¶0035-0040 and Fig. 4; and pgs. 16-17, ¶¶0041-43; pg. 26, ¶0057. Claim 16 further recites that the dynamic replication threshold is based, at least in part, on at least one dynamic measure of capacity. *Id.* pg. 17, ¶0043. And claim 16 further recites that “the requested object is served to the client from a server site other than the edge server site.” *Id.* pg. 10, ¶0028, Fig. 2, step 215, Fig. 3(a), step 345. *I.e.*, as recited in claim 16, the requested object is selectively copied to the edge server even though it is not served from the edge server.

Claim 23

As recited in independent claim 23 the invention described and claimed in the present application recites a computer program product embodied on non-transient computer-readable storage media including computer program code to cause a processor to perform a method for managed object replication and delivery in a system comprising a network having one or more edge server sites and one or more parent server sites distinct from said edge server sites. See *Specification*, Fig. 1 and pgs. 6-7, ¶¶0020-0022. The claimed computer program product includes computer code for: “directing a request by a client for an object to a particular edge server site in the network.” *Id.*, pg. 5, ¶0017 and pg. 8, ¶0024, at

² Here too the labels “[a]” and “[b]” are used here to merely to aid explanation, and are not intended to impose any limitations on the claim.

200 (in Fig. 2) and 300 (in Fig. 4).

The claimed computer program product also includes computer code for when “the particular edge server site does not have the requested object ... the particular edge server site redirecting the client request to a first server site distinct from the particular edge server site.” *Id.*, pgs. 9-10, ¶¶0026-0028 and Figs. 2 and 3 at 210, 310.

And the claimed computer program product further includes computer code for when “the particular edge server site does not have the requested object ..., then ... conditionally replicating a portion [pg. 26, ¶0057] of the requested object on the particular edge server site for use in serving future client requests [*Id.*, pg. 9, ¶¶0026-0027 and Figs. 2 and 3 at 315], said replicating being based at least in part on a dynamic measure of popularity of the requested object [*Id.*, pg. 12, ¶0032], wherein the portion of the requested object is replicated on the particular edge server site when

[X] the dynamic measure of popularity of the requested object
exceeds

[Y] a dynamic replication threshold [*Id.*, pgs. 16-17, ¶¶0041-43],

said dynamic replication threshold being based at least in part on at least one dynamic measure of capacity [*Id.* pg. 17, ¶0043], and wherein the requested object is served to the client from a server site other than the particular edge server site [*Id.* pg. 10, ¶0028, Fig. 2, step 215, Fig. 3(a), step 345].”³

Claim 38

As recited in independent claim 38 the invention described and claimed

³ Here too the labels “[X]” and “[Y]” are used here to merely to aid explanation, and are not intended to impose any limitations on the claim.

in the present application recites a “computer program product embodied on non-transient computer-readable storage media including computer program code.” As recited in the claim, the computer code has “instructions to cause a processor to perform a method for managed object replication and delivery in a system comprising a network having one or more parent server sites and one or more edge server sites distinct from said parent server sites.” *Specification*, Fig. 1 and pgs. 6-7, ¶¶0020-0022. The method includes directing a request by a client for an object to an optimal edge server site in the network [*Id.*, pg. 5, ¶0017 and pg. 8, ¶0024, at 200 (in Fig. 2) and 300 (in Fig. 4)]. The method further includes, if the optimal edge server site does not have the requested object, the edge server site redirecting the client request to a parent server site in the network distinct from said edge server site [*Id.*, pgs. 9-10, ¶¶0026-0028 and Figs. 2 and 3 at 210, 310]. The method further includes, if the edge server site does not have the requested object, conditionally replicating a portion of the requested object to the edge server site from a parent server site in the network, where the replicating is based at least in part on a dynamic measure of popularity of the requested object relative to a dynamic replication threshold. *Id.*, pg. 9, ¶¶0026-0027 and Figs. 2 and 3 at 315; pgs. 14-16, ¶¶0035-0040 and Fig. 4; and pgs. 16-17, ¶¶0041-43. As recited in the claim, “said dynamic replication threshold being based at least in part on at least a dynamic measure of capacity [*Id.* pg. 17, ¶0043]. Furthermore, if the optimal edge server site does not have the requested object then requested object is served to the client from a server site other than the edge server site – even if the object is replicated to the edge server site. *Id.* pg. 10, ¶0028, Fig. 2, step 215, Fig. 3(a), step 345.

Claim 45

As recited in independent claim 45 the invention described and claimed in the present application recites a system for managed object replication and

delivery. As recited in claim 45, the system includes a plurality of edge server sites in a network; and a plurality of parent server sites in the network, said parent server sites distinct from the edge server sites. *Specification*, Fig. 1 and pgs. 6-7, ¶¶0020-0022. As further recited in claim 45, “a request by a client for an object is directed to an edge server site in the network [*Id.*, pg. 5, ¶0017 and pg. 8, ¶0024, at 200 (in Fig. 2) and 300 (in Fig. 4)], and if the edge server site has the requested object, the requested object is served to the client from the edge server site.” The claim further recites that “if the edge server site does not have the requested object, the client request is redirected to a first server site distinct from the edge server site. *Id.*, pgs. 9-10, ¶¶0026-0028 and Figs. 2 and 3 at 210, 310. And the claim still further recites that “if the edge server site does not have the requested object and the object is served to the client from a server site other than the edge server site [*Id.*, pgs. 9-10, ¶¶0026-0028 and Figs. 2 and 3 at 210, 310], then a portion of the requested object is conditionally replicated on the edge server site for use in serving future client requests.” *Id.*, pg. 9, ¶¶0026-0027 and Figs. 2 and 3 at 315; pgs. 14-16, ¶¶0035-0040 and Fig. 4; and pgs. 16-17, ¶¶0041-43; pg. 26, ¶0057.

The claim further recites that the conditional replication on the edge server site is “based at least in part on *[A]* a dynamic measure of popularity of the requested object relative to *[B]* a dynamic replication threshold, said dynamic replication threshold being based, at least in part, on at least one dynamic measure of capacity.”⁴ *Id.*, pg. 9, ¶¶0026-0027 and Figs. 2 and 3 at 315; pgs. 14-16, ¶¶0035-0040 and Fig. 4; and pgs. 16-17, ¶¶0041-43.

⁴ Here too the labels “*[A]*” and “*[B]*” are used here to merely to aid explanation, and are not intended to impose any limitations on the claim.

**41.37(C)(1) (VI) GROUNDS OF REJECTION TO BE REVIEWED ON
APPEAL**

The ground of rejection presented for review on appeal is whether claims 1-3, 5-12, 14-18, 20-25, 27-34, 36-40, 42-47, 49-55, 57, 58, and 66-84 are unpatentable under 35 U.S.C. § 103(a) over Jungck and Sim.

41.37(C)(1)(VII) ARGUMENT

Each of the claims recites, in some form, selective replication of an object (or a portion of an object) to an edge server. The selective replication is made, based on *two* values, *viz.* (I) a dynamic *measure* of popularity of the object, and (II) **a dynamic threshold**. The Examiner has failed to show any teaching or suggestion in the prior art of any dynamic threshold, let alone of the use of such a dynamic threshold with a measure of popularity of an object.

The following table summarizes the selective/conditional replication recited in the claims:

Claims	Conditional action	First value (I)	Relationship	Dynamic threshold (II)
1 Dependents: 2-3, 5-12, 14-15, 66- 72, 82	conditionally replicating a portion of the requested object	dynamic measure of popularity of the requested object	exceeds	dynamic replication threshold ... based at least in part on at least one dynamic measure of capacity

Claims	Conditional action	First value (I)	Relationship	Dynamic threshold (II)
1 Dependents: 2-3, 5-12, 14-15, 66- 72, 82	conditionally replicating a portion of the requested object	dynamic measure of popularity of the requested object	exceeds	dynamic replication threshold ... based at least in part on at least one dynamic measure of capacity
16 Dependents: 17, 18, 20, 21, 22, 74- 75, 83	conditionally replicating a portion of the requested object to the edge server site from a parent server site in the network	dynamic measure of popularity of the requested object	relative to	dynamic replication threshold ... based, at least in part, on at least one dynamic measure of capacity
23 Dependents: 24, 25, 27- 34, 36-37, 73, 76-77	conditionally replicating a portion of the requested object on the particular edge server site for use in serving future client requests	dynamic measure of popularity of the requested object	exceeds	dynamic replication threshold ... based at least in part on at least one dynamic measure of capacity
38 Dependents: 39, 40, 42- 44, 78-79	conditionally replicating a portion of the requested object to the edge server site from a parent server site in the network,	dynamic measure of popularity of the requested object	relative to	dynamic replication threshold ... based at least in part on at least a dynamic measure of capacity

Claims	Conditional action	First value (I)	Relationship	Dynamic threshold (II)
1 Dependents: 2-3, 5-12, 14-15, 66- 72, 82	conditionally replicating a portion of the requested object	dynamic measure of popularity of the requested object	exceeds	dynamic replication threshold ... based at least in part on at least one dynamic measure of capacity
45 Dependents: 46, 47, 49- 55, 57, 58, 80-81, 84	a portion of the requested object is conditionally replicated on the edge server site for use in serving future client requests,	dynamic measure of popularity of the requested object	relative to	dynamic replication threshold ... based, at least in part, on at least one dynamic measure of capacity

Jungck and Sim, alone or in any proposed combination, do not teach or in any way suggest any *comparison of* (I) a *measure* of popularity of the object, and (II) a dynamic threshold, let alone any conditional replication based on such a comparison.

Claim 1 and its dependents are patentable over Jungck and Sim

Independent claim 1 recites, *inter alia*, “conditionally replicating a portion of the requested object on the particular edge server site for use in serving future client requests, said replicating being based at least in part on a dynamic measure of popularity of the requested object, wherein the portion of the requested object is replicated on the particular edge server site when

[i] the dynamic measure of popularity of the requested object
exceeds

[ii] a dynamic replication threshold.”⁵

However, there is no teaching or suggestion in either Jungck or Sim of any “dynamic replication threshold,” as required by the claim. And since there is no teaching of such a dynamic threshold in either Jungck or Sim, there is not and cannot be any teaching or suggestion of replication based on comparison of such “a dynamic replication threshold” with any other value, let alone with any “measure of popularity of the requested object.”

In **Jungck**, all client requests are handled by cache servers 208 (Jungck ¶0057), and *all* requested content is unconditionally cached. These cache servers intercept client requests and attempt to service them. *Id.* When a cache server does not have the requested content (a so-called “miss”), the cache forwards the request on to the content source which forwards the content to the client through the cache server (possibly acting as a proxy server). *Id.* “[T]he cache server 208 [then] saves a copy of the content in its cache for later requests.” *Id.* Unconditional caching has serious disadvantages discussed in the application as filed in Applicant’s earlier responses (caching large and infrequently requested objects is very inefficient).

Thus, as taught by Jungck, *all* requested content is (a) served by the cache server at which the client made the request, and (b) unconditionally cached. In both of these aspects, Jungck differs from the claimed invention (where (i) the requested content is served from a location (caching server site) other than that at which the request was made; and (ii) even though it is not being served from the location at which the client request was made, the content is conditionally cached at that location.)

The Examiner acknowledged that “Jungck did not explicitly disclose ... if the edge server site does not have the requested object, **conditionally replicating**

⁵ As noted above, the labels “[i]” and “[ii]” are used here to aid explanation.

a portion of the requested object to the edge server site from the parent server site in the network, wherein the portion of the requested object is replicated on the particular edge server site when the dynamic measure of popularity of the requested object exceeds a dynamic replication threshold.” *Final Office Action of August 3, 2011*, pg. 3. Applicant agrees, noting however, that Jungck did not *implicitly* disclose popularity-based replication either.

The Examiner states (at *Final Office Action of August 3, 2011*, pgs. 3-4) that Sim disclosed:

conditionally replicating {{Sim discloses that the content is replicated to the nodes of the network based on popularity of the content therefore there is a presence of "dynamic measurement" which determines content's popularity}} portion of the requested object to the edge server site from the parent server site in the network (paragraph 138), wherein the portion of the requested object is replicated on the particular edge server site when the dynamic measure of popularity of the requested object exceeds dynamic replication threshold (paragraphs 47, 52, 138, 197, 199) {when it is measured that popularity of the requested object has exceeded the threshold set for an object to be considered popular only then that object is replicated}, said dynamic replication threshold being based at least in part on at least one dynamic measure of capacity (Paragraphs. 230 and 236) {dynamic measure of capacity with respect to the dynamic replication of popular content is disclosed by Sim which specifically states that less popular content is removed to make room for more popular and new content when the available storage is running low. Here multiple dynamic aspects of invention are at play (i) identification of popular content (ii) Identification of the storage capacity of the storage space that hold the popular content (Hi) When the storage capacity is running low (I.E dynamic measure of capacity of the storage space in relation to the content being stored), comparison and identification of less popular content with more popular content and the removal of less popular to make room for more popular/relevant content in the storage space}}

Sim uses so-called distribution “stations” to distribute large payload files in a network. For example, Sim’s Figure 1 (reproduced here) shows stations denoted “STN – 1410”.

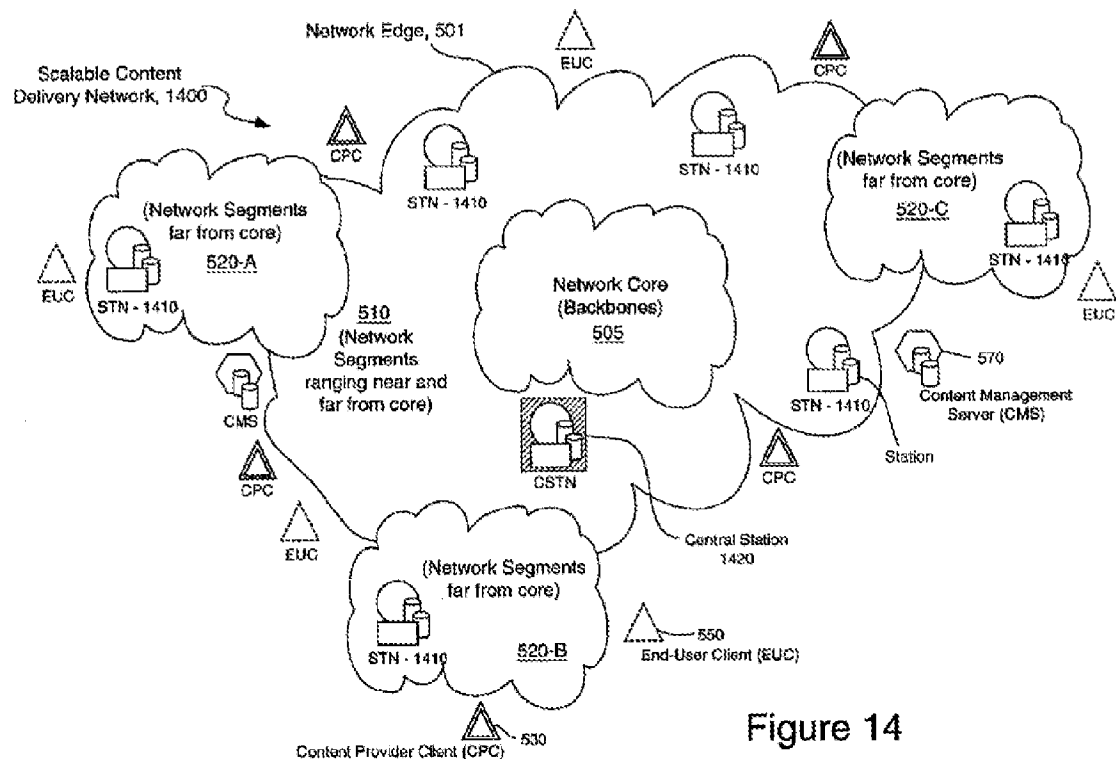


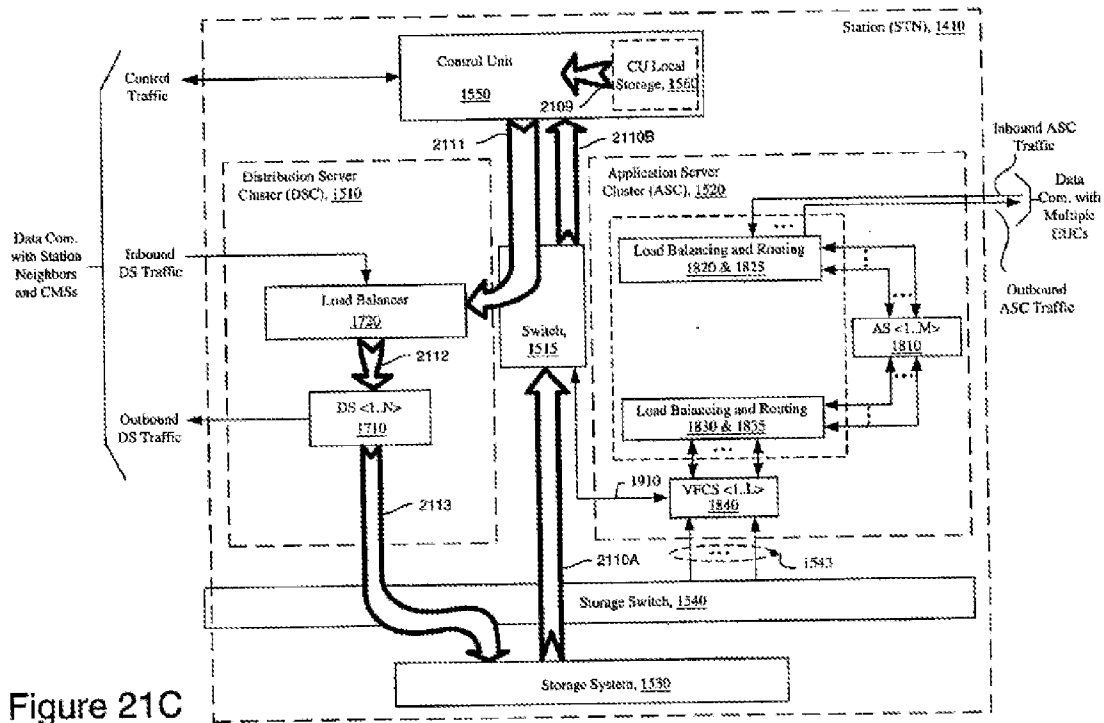
Figure 14

Sim’s figures. 21A through 21G “are illustrative examples of station operation and data flow.” Sim ¶¶0073.

Sim stores data about files in a File Metadata Database. *Id.* ¶¶0199, 0218, and ¶¶0224-0225. (“The File Metadata Database holds file metadata related to the block files which includes content provider ID, initial popularity index, block size, actual usage rating, media size, attributes, minimum retained size, access time stamps, track file indices, block file indices, storage volume indices, etc.”)

In Sim, “Storage Management Subsystem 2070 includes a Storage Management Agent as its primary component. The Storage Management Agent 1) monitors the total local storage availability and the storage availability of a content provider; 2) adjusts the “popularity” index of a file; 3) determines a storage safety

threshold for each content provider; 4) schedules content pruning; 5) computes pruning amount; 6) removes the least likely to be used blocks of a file; and 7) reports storage usage information and shortage warnings to Content Management servers.” *Sim*, ¶0197. These operations of the Storage Management Agent re described in greater detail in *Sim* with reference to Figure 21C (reproduced here below) which “illustrates retrieval of usage information, retrieval of content provider storage reservation and storage availability data, and related storage management activities commanded by the control unit and carried out by a selected DS.” *Sim* ¶0236.



As described in *Sim*’s Table 3, relied on by the Examiner:

In Step 2109, the Storage Management Subsystem 1560 within the Control Unit 1550 retrieves the file usage information. Then, in Step 2110A, the Storage Management Subsystem retrieves the Content Provider data, and reads the shared storage availability of the File Metadata Database. Next, in Step 2110B, the

Storage Management Subsystem uses the Content Provider data which it got from the File Metadata Database to identify the content provider's storage reservation. Then, in step 2111, Sim prunes data from the storage in order to make room for the next object (“based on the storage availability, content provider's storage reservation, media file usage information, and media file's storage occupancy, the Storage Management Subsystem instructs DS [the distribution server] ... to remove the least likely to be used blocks.”) *Id.*

“The Storage Management Subsystem also updates the file's popularity level.” *Id.* The “file” that Sim is referring to here is the file that is being copied into the local storage of the station, and, contrary to the Examiner’s assertion, the “popularity level” referred to in Sim ¶¶0236 is not a dynamic popularity threshold, it is the file’s popularity index.

Once the storage is pruned “an available DS ... removes blocks from the specified files and updates the File Metadata Database appropriately”. *Id.*

Sim describes use of popularity of a file (or of content) in a number of places (Abstract, ¶¶0025, 0047, 0052, 0105, 0138, 0197-199, 0225, 0230, 0236, and 0244), none of which in any way refer to or imply or suggest any kind of **dynamic threshold**.

Thus, while the cited portions of Sim (paragraphs 47, 52, 138, 197, and 199) do mention “popularity”, and adjusting a popularity index, and , Sim ¶¶230 and 236 discuss removal of less popular content from memory, those cited portions of Sim lack any mention of any dynamic threshold, as claimed. It appears that the Examiner has confused the *measure of popularity* with the *dynamic threshold* recited in the claim. The claim requires two values (denoted (i) and (ii) in the claim above), one of which is compared against the other. There is no teaching or suggestion of any dynamic threshold in Sim, and so there can be no comparison of such a value with anything else in Sim.

For at least these reasons, claim 1 and its dependents are patentable over Jungck and Sim.

Furthermore, claim 1 recites that “if the particular edge server site does not have the requested object, then ... conditionally replicating a portion of the requested object on the particular edge server site for use in serving future client requests, ... wherein *the requested object is served to the client from a server site other than the particular edge server site.*” As described in the application (and with reference to Fig. 2 of the application, reproduced here), the client request for an object is initially directed (at 200) to edge server 130 (an “optimal” edge server. *Specification*, ¶0024). “[I]f the selected edge server does not have the requested object, the selected edge server directs (at 210, 320) the requesting client 140 to a parent server 120.” *Id.* ¶0028. The parent server (or some other server) serves the requested object to the client (*e.g.*, 215 or 230). *Id.* ¶¶0028-0031. In the meantime, and separately from the requested object being served to the client, the originally selected edge server (130 in Fig. 2) selectively replicates the requested object. *Id.* ¶0032; Fig. 3(b).




Fig. 2

In other words, as recited in claim 1, the requested object is selectively copied to the edge server even though it is not served from the edge server.

The Examiner's rejection of all of the claims states: "Jungck did not explicitly disclose (C) if the edge server site does not have the requested object, conditionally replicating a portion of the requested object to the edge server site from the parent server site in the network, wherein the portion of the requested object is replicated on the particular edge server site when the dynamic measure of popularity of the requested object exceeds a dynamic replication threshold, said dynamic replication threshold being based at least in part on at least one dynamic measure of capacity." *Final Office Action of August 3, 2011*, pg. 3. But the Examiner has failed to address this element of the claims ("conditionally replicating a portion of the requested object on the particular edge server site for use in serving future client requests, ... wherein the requested object is served to the client from a server site other than the particular edge server site."). Sim

discloses distributing objects to nodes during an “initial file distribution phase” based on various factors, including “a ‘popularity’ index set by the content provider”. *Sim*, ¶0138 (cited by the Examiner). But this “initial file distribution phase” is not, as recited in claim 1 “conditionally replicating a portion of the requested object on the particular edge server site for use in serving future client requests, ... *wherein the requested object is served to the client from a server site other than the particular edge server site.*”

Sim describes and distinguishes four components or entities operating in his system. In *Sim*, an End-User Client (EUC) requests content from an Application Servers (AS) at which the request was made, and all content is served to the requesting client from the application server (AS). (*Sim* Fig. 5 and related description). Separate and apart from the content being served to the end users, *Sim* describes how Content Provider Client (CPC) systems 530 publish their content to the application servers 560 via one or more Content Management Servers (CMS) 570. (¶¶0082 *et seq.*) In *Sim*, the manner in which content is propagated from the content providers (CPC) to the content management servers is independent of the manner in which content is served.

In *Sim*, as with Jungck, *all* requested content is served by the server [Application Sever (AS) 560] at which the client made the request. However, in *Sim* content that is served is *not* cached at the server that serves that content. Instead, a so-called Virtual file control system (VCFS) server is used to make it appear to a requesting client that the content is present at the application server. (*Sim* ¶¶0047-48; see also “Each node at the edge of the network ... is configured to appear as if it has the large payload stored locally when portions of the file are really stored in on other nodes located throughout the network. ... When the end-user system issues a request for content (e.g., a large payload) the request is routed to the nearest node and the system delivers the requested content to the node in

manner that maximizes data transfer efficiency while minimizing bandwidth consumption. The end result is that each network node has access to numerous large data files without having to store each of those data files locally.” *Id.* ¶0080.)

More specifically, in Sim, the user always gets served from the same application server (AS) at which the request was made. (See ¶0047) If the application server does not have all of the requested data, that same application server obtains the needed data while it is downloading what it already has to the user. *Id.* (“since only selective portions of the large payload file may actually be resident at that node’s storage at the time of request, the distribution stations may download the non-resident portions of the file as the application server is servicing the user”).

Still further, Sim does not perform any conditional replication based on a measure of popularity of the requested object. Rather, as explained here, in Sim, popularity of already cached objects is used to determine which of those objects can be pruned from a cache to make room for new objects. This pruning is done independent of and in advance of any request for objects.

In Sim, large files may be replicated at nodes at the edge of the network. *Sim* ¶0047. However, this replication is made by a content provider in advance of any use of or request for the object, based, *e.g.*, on an initial *static* assessment of the objects popularity. *Id.* Thus, in *Sim*, if a content provider believes that certain objects will be popular, those are *unconditionally* replicated to the edge servers. This is akin to the prior art described in the present application (described, *e.g.*, at ¶0013, which notes that “it is difficult to predict popularity and difficult to manage pre-populating.”). *Sim* then uses popularity elsewhere, also *not to decide what to*

cache, but rather for cache pruning – *i.e.*, to decide what to *remove* from a cache.

Sim, ¶0236

Sim provides a so-called Storage Management Agent (SMA) to administer aspects of the storage system. The operation of this SMA is separate and apart from Sim’s serving and associated replication process. As Sim explains, one function of the SMA is to try to maintain sufficient free space on the various storage systems.

The Storage Management Agent ... determine[s] a reasonable storage safety threshold, adjusts the "popularity" index of a file, and identifies the least likely to be used blocks. A storage safety threshold is the minimum amount of free storage each content provider must reserve at all times. Based on storage availability and the DS activities, the Storage Management Agent determines the total amount of data to be pruned for each content provider and schedules the deletion of the least likely to be used blocks.

Sim ¶0199, emphasis provided.

Sim thus simply uses a “popularity index” as part of a management process to try to *maintain* sufficient free space on a storage system and so to decide *what to prune* in a storage system (“... decides what content to prune ...” *Sim*, Abstract). The “popularity” index mentioned in Sim is not used to decide whether or not to replicate an object that is being served to a user (in Sim, as in Jungck, all replication is unconditional). As Sim states at ¶0052 (with emphasis added):

... the portions and amount of a large payload file ***maintained*** at each node depends on the available storage, popularity of the content, distribution criteria by the content provider, etc. Thus, least likely to be used blocks of a large payload file ***may be pruned*** (*i.e.*, deleted from local storage) *to make room for other highly desirable content*.

See also *Sim* at ¶0230 which states:

The Storage Management Subsystem watches the available shared storage, the content provider's reserved storage, and the usage logs. It initiates the removal of less popular content to make room for more popular and new content when available storage is running low.

At the time a cache is pruned in Sim, there is no requested object, and so, in Sim, to the extent popularity is used, it is not the popularity of the requested object. Rather, it is the popularity of other cached objects that are used to determine which of them to prune (remove) from the cache.

Furthermore, in any such combination, all content would be served from the server at which the request was originally made. This feature is contrary to the claimed requirement that “the requested object is served to the client from a server site other than the particular edge server site.”

For at least these additional reasons, claim 1 is patentable over Jungck and Sim.

Claims 2, 3, 5-12, 14, 15, and 66-72 depend from claim 1 and are therefore patentable over Sim and Jungck for at least the reasons given above.

Claim 72 (depends from claim 1) recites “the method of claim 1 wherein the requested object comprises chunks, including initial chunks and remaining chunks, and wherein the portion of the requested object comprises only initial chunks of the object.” This feature is neither taught nor suggested by Jungck or Sim. The Examiner relies on Sim paragraph 138 to teach this claimed replication of “only initial chunks of the object.” *Final Office Action of August 3, 2011* §20, pg. 13.

Sim paragraph 129 teaches that a movie (or the like) can be treated as a series of blocks. Sim also teaches that the blocks can be accessed “out of order”. *Id.* And Sim also teaches that “pruning (freeing the storage used by some blocks for use by other blocks) can be done at the ‘block level’”. *Id.* Sim paragraph 144

relates to pruning. But nowhere does Sim teach or in any way suggest any policy for *replication* at the block level, let alone for a policy as recited in claim 72 that “replicates only initial chunks of the object.”

For at least these additional reasons claim 72 is further patentable over Jungck and Sim.

Claim 16 and its dependents are patentable over Jungck and Sim

Claim 16 recites, *inter alia*, “conditionally replicating a portion of the requested object to the edge server site from a parent server site in the network, said replicating being based at least in part on

[a] a dynamic measure of popularity of the requested object
relative to

[b] a dynamic replication threshold,

said dynamic replication threshold being based, at least in part, on at least one dynamic measure of capacity”

Here, again, there is no teaching or suggestion in either Jungck or Sim of any dynamic threshold, as required by the claim. The Examiner relies on Sim to teach these features of claim 16, repeating the same argument he used for claim 1 and the other independent claims.

Here again (as with claim 1), it appears that the Examiner has confused the measure of popularity with the “dynamic replication threshold” recited in the claim. The claim requires replication based at least in part on a relationship between two values (denoted (a) and (b) in the claim above). There is no teaching or suggestion of “dynamic replication threshold” in Sim, and so there can be no comparison of such a value with anything else in Sim.

Furthermore, as with claim 1, the Examiner has ignored an element of the claim (“if the edge server site does not have the requested object, conditionally replicating a portion of the requested object to the edge server site from a parent server site in the network, ... *wherein the requested object is served to the client from a server site other than the edge server site.*”) As noted above, Sim discloses distributing objects during an “initial file distribution phase” based on various factors, including “a ‘popularity’ index set by the content provider”. *Sim*, ¶0138 (cited by the Examiner). But Sim’s “initial file distribution phase” is not, as recited in claim 16 “conditionally replicating a portion of the requested object to the edge server site from a parent server site in the network, ... *wherein the requested object is served to the client from a server site other than the edge server site.*”

For at least these reasons, claim 16 and its dependents are patentable over Jungck and Sim.

Claim 23 and its dependents are patentable over Jungck and Sim

Claim 23 recites, *inter alia*, “conditionally replicating a portion of the requested object on the particular edge server site for use in serving future client requests, said replicating being based at least in part on a dynamic measure of popularity of the requested object, wherein the portion of the requested object is replicated on the particular edge server site when

[p] the dynamic measure of popularity of the requested object
exceeds

[q] a dynamic replication threshold, said dynamic replication threshold
being based at least in part on at least one dynamic measure of capacity.”

The Examiner combines the rejection of claim 23 over Jungck and Sim

with the rejections of claims 1, 16, 38, and 45 over those references.

Here again (as with claims 1 and 16), it appears that the Examiner has confused the “dynamic measure of popularity of the requested object” with the “dynamic replication threshold” recited in the claim. The claim requires two values (denoted (p) and (q) in the claim above), one of which is compared against the other.

Here, again, there is no teaching or suggestion in either Jungck or Sim of any dynamic popularity threshold, as required by the claim. And since there is no teaching of such a threshold in either Jungck or Sim, there is not and cannot be any teaching or suggestion of replication based on comparison of such a dynamic threshold with any other value, let alone with any “dynamic measure of popularity of the requested object.”

The Examiner again relies on Sim to teach these features of claim 23, combining the rejection of claim 23 with those of claims 1, 16, 38, and 45.

And here, again, the Examiner has ignored a claim element (“if the particular edge server site does not have the requested object, then ... conditionally replicating a portion of the requested object on the particular edge server site for use in serving future client requests, ... wherein the requested object is served to the client from a server site other than the particular edge server site.”)

For the same reasons given above for claims 1 and 16, claim 23 and its dependents are patentable over Jungck and Sim.

Claim 38 and its dependents are patentable over Jungck and Sim

Claim 38 recites, *inter alia*, “conditionally replicating a portion of the requested object to the edge server site from a parent server site in the network, said replicating being based at least in part on a dynamic measure of popularity of

the requested object relative to a dynamic replication threshold, said dynamic replication threshold being based at least in part on at least a dynamic measure of capacity.” Claim 38 also recites “if the edge server site does not have the requested object, conditionally replicating a portion of the requested object to the edge server site from a parent server site in the network ... wherein the requested object is served to the client from a server site other than the edge server site.”

The Examiner again relies on Sim to teach these features of claim 38. For the same reasons given above for claims 1 and 16, claim 38 and its dependents are patentable over Jungck and Sim.

Claim 45 and its dependents are patentable over Jungck and Sim

Claim 45 recites, *inter alia*, “if the edge server site does not have the requested object and the object is served to the client from a server site other than the edge server site, then a portion of the requested object is conditionally replicated on the edge server site for use in serving future client requests, said replicating being based at least in part on a $[x]$ /dynamic measure of popularity of the requested object relative to $[y]$ a dynamic replication threshold, said dynamic replication threshold being based, at least in part, on at least one dynamic measure of capacity.”

In addition to Sim’s lack of any dynamic replication threshold, Sim lacks any teaching or suggestion of the claimed “if the edge server site does not have the requested object and the object is served to the client from a server site other than the edge server site, then a portion of the requested object is conditionally replicated on the edge server site for use in serving future client requests.”

The Examiner again relies on Sim to teach these features of claim 45, and for at least the reasons given above with respect to claims 1, 16, 23, and 38, claim 45 and its dependents are patentable over Jungck and Sim.

41.37(C)(1) (VIII) CLAIMS APPENDIX

Claim 1. A method for managed object replication and delivery in a system comprising a network having one or more edge server sites and one or more parent server sites distinct from said edge server sites, the method comprising:

- (A) directing a request by a client for an object to a particular edge server site in the network; and
- (B) if the particular edge server site does not have the requested object, then
 - (b1) said particular edge server site redirecting the client request to a first server site distinct from the edge server site; and
 - (b2) conditionally replicating a portion of the requested object on the particular edge server site for use in serving future client requests, said replicating being based at least in part on a dynamic measure of popularity of the requested object, wherein the portion of the requested object is replicated on the particular edge server site when the dynamic measure of popularity of the requested object exceeds a dynamic replication threshold, said dynamic replication threshold being based at least in part on at least one dynamic measure of capacity, and wherein the requested object is served to the client from a server site other than the particular edge server site.

Claim 2. The method of claim 1, wherein redirecting the client request to said first server site comprises said particular edge server site redirecting the client request to a parent server site in the network.

Claim 3. The method of claim 2, wherein if that parent server site does not have the requested object, then recursively redirecting the request until a parent server site in the network having the requested object is reached, and then serving the requested object to the client from the parent server site that has the requested object.

Claim 4. (Canceled)

Claim 5. The method of claim 1, wherein directing a request by a client for an object to a particular edge server site comprises directing the request by the client for an object to a best or optimal edge server site.

Claim 6. The method of claim 5, wherein a best or optimal edge server site comprises an edge server site selected using at least one of a determination based on a best repeater selector, the likelihood of a copy of the requested object being available at the edge server site, and the bandwidth between the edge server site and the client.

Claim 7. The method of claim 1, wherein said step of conditionally replicating the portion of the requested object to the particular edge server site comprises replicating the portion of the requested object to the particular edge server site from a parent server site.

Claim 8. The method of claim 1, wherein said step of conditionally replicating comprises:

if the requested object is determined to be popular based on said dynamic measure of popularity, and if the requested object is unavailable on parent server sites in the network, then replicating the portion of the requested object to a parent server site in the network from an origin server site.

Claim 9. The method of claim 1, wherein said step of conditionally replicating the portion of the requested object to the particular edge server site comprises, if the requested object is unavailable on parent server sites in the network, replicating the portion of the requested object to the particular edge server site from an origin server site.

Claim 10. The method of claim 1, wherein said dynamic measure of popularity of the requested object is determined using at least a request rate for the requested object.

Claim 11. The method of claim 1, further comprising:

if an object on the particular edge server site is no longer popular, as determined based on said dynamic measure of popularity of said object, deleting at least some part of the object from the particular edge server site.

Claim 12. The method of claim 1, further comprising:

if an object on the parent server site is no longer popular, as determined based on said dynamic measure of popularity of said object, and if the object is available on an origin server site, then deleting at least some part of the object from the parent server site.

Claim 13. (Canceled).

Claim 14. The method of claim 1, wherein said step of conditionally replicating the portion of the requested object on said particular edge server site comprises:

replicating the portion of the requested object when said dynamic measure of popularity of the requested object is greater than the dynamic replication threshold and there is enough storage on said particular edge server site to replicate the portion of the requested object;

otherwise, if there is not enough storage on said particular edge server site to replicate the portion of the requested object, then

i) comparing the dynamic measure of popularity of the requested object against a dynamic measure of popularity of a least popular object in storage on said particular edge server site,

ii) if the dynamic measure of popularity of the requested object exceeds the dynamic measure of popularity of the least popular object in the storage, deleting at least some part of the least popular object from the storage, and then

iii) repeating i) and ii) until enough storage is available for the portion of the requested object or until the dynamic measure of popularity of the requested object is less than the dynamic measure of popularity of the least popular object in the storage, and then

iv) replicating the portion of the requested object on said particular edge server site if there is enough storage on said particular edge server site.

Claim 15. The method of claim 1, wherein the step of serving the requested object is performed separately from the step of conditionally replicating the portion of the requested object.

Claim 16. A method for managed object replication and delivery in a system comprising a network having one or more parent server sites and one or more edge server sites distinct from said parent server sites, the method comprising:

(A) directing a request by a client for an object to an optimal edge server site in the network; and

(B) if the edge server site does not have the requested object, then the edge server site redirecting the client request to a parent server site in the network distinct from the edge server site; and,

(C) if the edge server site does not have the requested object, conditionally replicating a portion of the requested object to the edge server site from a parent server site in the network, said replicating being based at least in part on a dynamic measure of popularity of the requested object relative to a dynamic replication threshold, said dynamic replication threshold being based, at least in part, on at least one dynamic measure of capacity,

wherein the requested object is served to the client from a server site other than the edge server site.

Claim 17. The method of claim 16, further comprising:

if an object on the edge server site is no longer popular, as determined based on said dynamic measure of popularity of said object, deleting at least some

part of the object from the edge server site.

Claim 18. The method of claim 16, further comprising:

if an object on the parent server site is no longer popular, as determined based on the said dynamic measure of popularity of said object, and if the object is available on an origin server site, deleting at least some part of the object from the parent server site.

Claim 19. (Canceled)

Claim 20. The method of claim 16, wherein said step of conditionally replicating the portion of the requested object on said edge server site comprises:

replicating the portion of the requested object when said dynamic measure of popularity of the requested object is greater than the dynamic replication threshold and there is enough storage on said edge server site to replicate the portion of the requested object;

otherwise, if there is not enough storage on said edge server site to replicate the portion of the requested object, then

i) comparing the popularity of the requested object, as determined based on said dynamic measure of popularity, against a dynamic measure of popularity of a least popular object in storage on the edge server site,

ii) if the dynamic measure of popularity of the requested object exceeds the dynamic measure of popularity of the least popular object in the storage, deleting at least some part of the least popular object from the storage, and

iii) repeating i) and ii) until enough storage is available for the portion of the requested object or until the dynamic measure of popularity of the requested object is less than the dynamic measure of popularity of the least popular object in the storage, and then

iv) replicating the portion of the requested object on the edge server if there is enough storage.

Claim 21. The method of claim 16, wherein whether the requested object is popular is determined using at least a request rate for the requested object.

Claim 22. The method of claim 16, wherein the step of serving the requested object is performed separately from the step of conditionally replicating a portion of the requested object.

Claim 23. A computer program product embodied on non-transient computer-readable storage media including computer program code to cause a processor to perform a method for managed object replication and delivery in a system comprising a network having one or more edge server sites and one or more parent server sites distinct from said edge server sites, the computer-readable media comprising computer program code for:

(A) directing a request by a client for an object to a particular edge server site in the network; and

(B) if the particular edge server site does not have the requested object, then

(b1) the particular edge server site redirecting the client request to a

first server site distinct from the particular edge server site; and,

(b2) conditionally replicating a portion of the requested object on the particular edge server site for use in serving future client requests, said replicating being based at least in part on a dynamic measure of popularity of the requested object, wherein the portion of the requested object is replicated on the particular edge server site when the dynamic measure of popularity of the requested object exceeds a dynamic replication threshold, said dynamic replication threshold being based at least in part on at least one dynamic measure of capacity, and wherein the requested object is served to the client from a server site other than the particular edge server site.

Claim 24. The computer program product of claim 23, wherein redirecting the client request to a first server site comprises redirecting the client request to a parent server site in the network.

Claim 25. The computer program product of claim 23, wherein redirecting the client request to a first server site comprises redirecting the client request to a parent server site in the network, and if that parent server site does not have the requested object, then recursively redirecting the request until a parent server site in the network having the requested object is reached and then serving the requested object to the client from the parent server site that has the requested object.

Claim 26. (Canceled)

Claim 27. The computer program product of claim 23, wherein directing a request by a client for an object to a particular edge server site comprises directing the request by the client for an object to a best or optimal edge server site.

Claim 28. The computer program product of claim 27, wherein a best or optimal edge server site comprises an edge server site selected using at least one of a determination based on a best repeater selector, the likelihood of a copy of the requested object being available at the edge server site, and the bandwidth between the edge server site and the client.

Claim 29. The computer program product of claim 23, wherein the conditionally replicating the portion of the requested object to the particular edge server site comprises replicating the portion of the requested object to the particular edge server site from a parent server site.

Claim 30. The computer program product of claim 23, the method further comprising:

if the requested object is determined to be popular based on said dynamic measure of popularity, and if the requested object is unavailable on parent server sites in the network, then replicating the portion of the requested object to a parent server site in the network from an origin server site.

Claim 31. The computer program product of claim 23, wherein said conditionally replicating the portion of the requested object to the particular edge

server site comprises, if the requested object is unavailable on parent server sites in the network, replicating the portion of the requested object to the particular edge server site from an origin server site.

Claim 32. The computer program product of claim 23, wherein said dynamic measure of popularity of the requested object is determined using at least a request rate for the requested object.

Claim 33. The computer program product of claim 23, the method further comprising:

if an object on the particular edge server site is no longer popular, as determined based on said dynamic measure of popularity of said object, deleting at least some part of the object from the particular edge server site.

Claim 34. The computer program product of claim 23, the method further comprising:

if an object on the parent server site is no longer popular, as determined based on said dynamic measure of popularity of said object, and if the object is available on an origin server site, then deleting at least some part of the object from the parent server site.

Claim 35. (Canceled)

Claim 36. The computer program product of claim 23, wherein said

conditionally replicating the portion of the requested object on said particular edge server site comprises:

replicating the portion of the requested object when said dynamic measure of popularity of the requested object is greater than the dynamic replication threshold and there is enough storage on said particular edge server site to replicate the portion of the requested object;

otherwise, if there is not enough storage on said particular edge server site to replicate the portion of the requested object,

i) comparing the dynamic measure of popularity of the requested object against a dynamic measure of popularity of a least popular object in storage on said particular edge server site,

ii) if the dynamic measure of popularity of the requested object exceeds the dynamic measure of popularity of the least popular object in the storage, deleting at least some part of the least popular object from the storage, and

iii) repeating i) and ii) until enough storage is available for the portion of the requested object or until the dynamic measure of popularity of the requested object is less than the dynamic measure of popularity of the least popular object in the storage, and then

iv) replicating the portion of the requested object on said particular edge server site if there is enough storage.

Claim 37. The computer program product of claim 23, wherein serving the requested object is performed separately from conditionally replicating the portion of the requested object.

Claim 38. A computer program product embodied on non-transient computer-readable storage media including computer program code having instructions to cause a processor to perform a method for managed object replication and delivery in a system comprising a network having one or more parent server sites and one or more edge server sites distinct from said parent server sites, the method comprising:

(A) directing a request by a client for an object to an optimal edge server site in the network; and

(B) if the optimal edge server site does not have the requested object, the edge server site redirecting the client request to a parent server site in the network distinct from said edge server site; and,

(C) if the edge server site does not have the requested object, conditionally replicating a portion of the requested object to the edge server site from a parent server site in the network, said replicating being based at least in part on a dynamic measure of popularity of the requested object relative to a dynamic replication threshold, said dynamic replication threshold being based at least in part on at least a dynamic measure of capacity, wherein the requested object is served to the client from a server site other than the edge server site.

Claim 39. The computer program product of claim 38, the method further comprising:

if an object on the edge server site is no longer popular, as determined based on said dynamic measure of popularity of said object, deleting at least some part of the object from the edge server site.

Claim 40. The computer program product of claim 38, the method further comprising:

if an object on the parent server site is no longer popular as determined based on said dynamic measure of popularity of said object, and if the object is available on an origin server site, deleting at least some part of the object from the parent server site.

Claim 41. (Canceled)

Claim 42. The computer program product of claim 38, wherein replicating the portion of the requested object on said edge server site comprises:

replicating the portion of the requested object when said dynamic measure of popularity of the requested object is greater than the dynamic replication threshold and there is enough storage on said edge server site to replicate the portion of the requested object;

otherwise, if there is not enough storage on said edge server site to replicate the portion of the requested object, then

i) comparing the popularity of the requested object, as determined based on said dynamic measure of popularity, against a dynamic measure of popularity of a least popular object in storage on the edge server site,

ii) if the dynamic measure of popularity of the requested object exceeds the dynamic measure of popularity of the least popular object in the storage, deleting at least some part of the least popular object from the storage, and

iii) repeating i) and ii) until enough storage is available for the portion of the requested object or until the dynamic measure of popularity of the requested

object is less than the dynamic measure of popularity of the least popular object in the storage, and then

iv) replicating the portion of the requested object on the edge server if there is enough storage.

Claim 43. The computer program product of claim 38, wherein whether the requested object is popular is determined using at least a request rate for the requested object.

Claim 44. The computer program product of claim 38, wherein serving the requested object is performed separately from replicating the portion of the requested object.

Claim 45. A system for managed object replication and delivery, comprising:

a plurality of edge server sites in a network; and

a plurality of parent server sites in the network, said parent server sites distinct from said edge server sites,

wherein:

a request by a client for an object is directed to an edge server site in the network, and

if the edge server site has the requested object, the requested object is served to the client from the edge server site, and

otherwise, if the edge server site does not have the requested object, the

client request is redirected to a first server site distinct from the edge server site, and,

if the edge server site does not have the requested object and the object is served to the client from a server site other than the edge server site, then a portion of the requested object is conditionally replicated on the edge server site for use in serving future client requests, said replicating being based at least in part on a dynamic measure of popularity of the requested object relative to a dynamic replication threshold, said dynamic replication threshold being based, at least in part, on at least one dynamic measure of capacity.

Claim 46. The system of claim 45, wherein redirecting the client request to a server site comprises redirecting the client request to a parent server site in the network and attempting to serve the requested object to the client from the parent server site.

Claim 47. The system of claim 45, wherein redirecting the client request to a server site comprises recursively redirecting the request until a parent server site in the network having the requested object is reached and serving the requested object to the client from the parent server site.

Claim 48. (canceled)

Claim 49. The system of claim 45, wherein directing a request by a client for an object to an edge server site comprises directing the request by the client for an object to a best or optimal edge server site.

Claim 50. The system of claim 49, wherein a best or optimal edge server site comprises an edge server site selected using at least one of a determination based on a best repeater selector, the likelihood of a copy of the requested object being available at the edge server site, and the bandwidth between the edge server site and the client.

Claim 51. The system of claim 45, wherein replicating of a portion of the requested object to the edge server site comprises replicating the portion of the requested object to the edge server site from a parent server site.

Claim 52. The system of claim 45, wherein at least one of the plurality of edge server sites and the plurality of parent server sites further replicate the portion of the requested object from an origin server site if the requested object is popular, as determined based on the dynamic measure of popularity, and if the requested object is unavailable on parent server sites in the network.

Claim 53. The system of claim 45, wherein replicating the portion of the requested object to the edge server site comprises, if the requested object is unavailable on parent server sites in the network, replicating the portion of the requested object to the edge server site from an origin server site.

Claim 54. The system of claim 45, wherein said dynamic measure of popularity of the requested object is determined using at least a request rate for the requested object.

Claim 55. The system of claim 45, wherein at least one of the plurality of edge server sites and the plurality of parent server sites deletes at least some part of an object if the object is no longer popular, as determined based on said dynamic measure of popularity of the requested object.

Claim 56. (Canceled)

Claim 57. The system of claim 45, wherein replicating the portion of the requested object comprises:

replicating the portion of the requested object when said dynamic measure of popularity of the requested object is greater than the dynamic replication threshold and when there is enough storage on said the edge server site to replicate the portion of the requested object;

otherwise, if there is not enough storage on the edge server site to replicate the portion of the requested object, then

i) comparing the popularity of the requested object, as determined based on said dynamic measure of popularity of said object, against a dynamic measure of popularity of a least popular object in storage on the edge server site,

ii) if the dynamic measure of popularity of the requested object exceeds the dynamic measure of popularity of the least popular object in the storage, deleting at least some part of the least popular object from the storage, and

iii) repeating i) and ii) until enough storage is available for the portion of the requested object or until the dynamic measure of popularity of the requested object is less than the dynamic measure of popularity of a least popular object in

the storage, and then

iv) replicating the portion of the requested object on said edge server site if there is enough storage.

Claim 58. The system of claim 45, wherein serving the requested object is performed separately from replicating the portion of the requested object.

Claims 59- 65. (Canceled).

Claim 66. The method of claim 1 wherein the server site from which the requested object is served to the client is a peer server site of the particular edge server site.

Claim 67. The method of claim 1 wherein the server site from which the requested object is served to the client is the first server site.

Claim 68. The method of claim 1 wherein the server site from which the requested object is served to the client is a peer of the first server site.

Claim 69. The method of claim 1 wherein the step of conditionally replicating the portion of the requested object on the particular edge server site replicates the portion of the requested object from a peer server site of the particular edge server site.

Claim 70. The method of claim 1 wherein the step of conditionally replicating the portion of the requested object on the particular edge server site replicates the portion of the requested object from a server site.

Claim 71. The method of claim 1 wherein the dynamic measure of popularity of the requested object is based at least in part on one or more of:

- (a) a local dynamic measure of popularity of the object; and
- (b) information regarding the popularity of the object on other servers.

Claim 72. The method of claim 1 wherein the requested object comprises chunks, including initial chunks and remaining chunks, and wherein the portion of the requested object comprises only initial chunks of the object.

Claim 73. The method of claim 27 wherein the remaining chunks of the object are selectively replicated based at least in part on how many of the initial chunks are served to the client.

Claim 74. The method of claim 16 wherein the dynamic measure of popularity of the requested object is based at least in part on one or more of:

- (a) a local dynamic measure of popularity of the object; and
- (b) information regarding the popularity of the object on other servers.

Claim 75. The method of claim 16 wherein the requested object comprises chunks, including initial chunks and remaining chunks, and wherein the portion of

the requested object comprises only initial chunks of the object.

Claim 76. The computer program product of claim 23 wherein the dynamic measure of popularity of the requested object is based at least in part on one or more of:

- (a) a local dynamic measure of popularity of the object; and
- (b) information regarding the popularity of the object on other servers.

Claim 77. The computer program product of claim 23 wherein the requested object comprises chunks, including initial chunks and remaining chunks, and wherein the portion of the requested object comprises only initial chunks of the object.

Claim 78. The computer program product of claim 38 wherein the dynamic measure of popularity of the requested object is based at least in part on one or more of:

- (a) a local dynamic measure of popularity of the object; and
- (b) information regarding the popularity of the object on other servers.

Claim 79. The computer program product of claim 38 wherein the requested object comprises chunks, including initial chunks and remaining chunks, and wherein the portion of the requested object comprises only initial chunks of the object.

Claim 80. The system of claim 45 wherein the dynamic measure of popularity of the requested object is based at least in part on one or more of:

- (a) a local dynamic measure of popularity of the object; and
- (b) information regarding the popularity of the object on other servers.

Claim 81. The system of claim 45 wherein the requested object comprises chunks, including initial chunks and remaining chunks, and wherein the portion of the requested object comprises only initial chunks of the object.

Claim 82. A method as recited in claim 1, wherein the at least one dynamic measure of capacity represents available capacity on the particular edge server.

Claim 83. A method as recited in claim 16, wherein the at least one dynamic measure of capacity represents available capacity on the edge server site.

Claim 84. A system as recited in claim 45, wherein the at least one dynamic measure of capacity represents available capacity on the edge server site.

In re Application of: SEED, Steven, et al.
Application S.N.: 10/073,938
Appeal Brief

41.37(C)(1)(IX) EVIDENCE APPENDIX

There are no documents or items being submitted under this section.

41.37(C)(1)(X) RELATED PROCEEDINGS APPENDIX

There are no documents or items being submitted under this section.

CONCLUSION

Applicant requests on the basis of the arguments presented above that the outstanding rounds for rejection be reversed. Applicant submits that the application is in condition for allowance.

Respectfully submitted,

/David. D. Wier /

Date: November 29, 2011

David D. Wier, Attorney Reg. No. 48,229
Level 3 Communications, LLC
1025 Eldorado Blvd.
Broomfield, CO 80021
(720) 888-3539